

A universal approximate cross-validation criterion and its asymptotic distribution

Daniel Commenges^{1,2}

Cécile Proust-Lima^{1,2}

Cécilia Samieri^{1,2}

Benoit Liqueur^{1,2}

March 7, 2013

¹ INSERM, ISPED, Centre INSERM U-897-Epidemiologie-Biostatistique, Bordeaux, F-33000

² Univ. Bordeaux, ISPED, Centre INSERM U-897-Epidemiologie-Biostatistique, Bordeaux, F-33000, France

ABSTRACT. A general framework is that the estimators of a distribution are obtained by minimizing a function (the estimating function) and they are assessed through another function (the assessment function). The estimating and assessment functions generally estimate risks. A classical case is that both functions estimate an information risk (specifically cross entropy); in that case Akaike information criterion (AIC) is relevant. In more general cases, the assessment risk can be estimated by leave-one-out crossvalidation. Since leave-one-out crossvalidation is computationally very demanding, an approximation formula can be very useful. A universal approximate crossvalidation criterion (UACV) for the leave-one-out crossvalidation is given. This criterion can be adapted to different types of estimators, including penalized likelihood and maximum a posteriori estimators, and of assessment risk functions, including information risk functions and continuous rank probability score (CRPS). This formula reduces to Takeuchi information criterion (TIC) when cross entropy is the risk for both estimation and assessment. The asymptotic distribution of UACV and of a difference of UACV is given. UACV can be used for comparing estimators of the distributions of ordered categorical data derived from threshold models and models based on continuous approximations. A simulation study and an analysis of real psychometric data are presented.

Key Words: AIC; cross entropy; crossvalidation; estimator choice; Kullback-Leibler risk; model selection; ordered categorical observations; psychometric tests

1 Introduction

Akaike information criterion (AIC) (Akaike, 1973) is useful for comparing parametric models. AIC assumes parametric models and maximum likelihood estimators; it has been developed from the Kullback-Leibler divergence. Commenges et al. (2008) showed that a normalized difference of AIC of two models could be considered as estimating a difference of Kullback-Leibler risks of maximum likelihood estimators of the density based on the two models. Likelihood crossvalidation (LCV) has also been widely used for comparing parametric models. Stone (1974) showed that LCV was asymptotically identical to AIC. LCV however is more flexible in that it can be applied to other estimators than maximum likelihood estimators (MLE), for instance to penalized likelihood estimators: see Golub et al. (1979); Wahba (1985). It has asymptotic optimality properties (Van Der Laan et al., 2004). The leave-one-out crossvalidation is the most natural and one of the most efficient but it is also the most computationally demanding so that approximation formulas have been derived. Stone (1974) was the first to give an approximation formula. Other works developed generalized approximation crossvalidation for smoothing parameter selection for penalized splines (Xiang and Wahba, 1996; Gu and Xiang, 2001) or for penalized likelihood (O’Sullivan, 1986; Commenges et al., 2007). Cross-validation can also be applied to other assessment risks than Kullback-Leibler risk; for a review see Arlot and Celisse (2010).

We consider the following framework: estimators of the true density function are defined as minimizing an estimating function; the estimating function itself can be viewed as an estimator of a risk, that we call an “estimating risk”; the estimators of the true density are assessed using an “assessment risk”. The assessment risk can be estimated by crossvalidation, allowing a choice between them. Leave-one-out crossvalidation is one of the best crossvalidation procedure but is very computationally demanding. The aim of this paper is to find a universal approximation for leave-one-out crossvalidation, valid whatever the estimating and assessment risks.

Section 2 presents the framework and the universal approximate criterion (UACV); section 3 shows how it specializes to particular cases. In section 4 the asymptotic distributions of UACV and of a difference of UACV are given. Section 5 presents a simulation study, while section 6 presents an illustration of the use of UACV for comparing estimators derived from threshold models and estimators obtained by continuous approximations in the case of ordered categorical data with repeated measurements. Section 7 concludes.

2 The universal approximate crossvalidation criterion

2.1 The risk and its estimation by crossvalidation

Suppose that a sample of independently identically distributed (iid) variables $\bar{O}_n = (Y_i, i = 1, \dots, n)$ is available. Based on \bar{O}_n , an estimator $g^{\hat{\theta}}$ (where $\hat{\theta}$ is short for $\hat{\theta}_n$) of the probability density function f^* of the true distribution can be chosen in a family of distributions $(g^\theta)_{\theta \in \Theta}$, $\Theta \subset \mathbb{R}^p$. Estimators are chosen as minimizing an estimating function $\Phi_{\bar{O}_n}(\theta)$ where $\Phi_{\bar{O}_n}(\theta) = n^{-1} \sum_{i=1}^n \phi(\theta, Y_i)$: that is, $\hat{\theta} = \operatorname{argmin}_{\theta} \Phi_{\bar{O}_n}(\theta)$. We shall assume that $\phi(\theta, y)$ is a thrice differentiable function of θ for any y . Under mild conditions $\Phi_{\bar{O}_n}$ converges toward $\Phi_\infty(\theta) = E_*\{\phi(\theta, Y_i)\}$, where E_* means that the expectation is taken under the true probability specified by f^* . Thus, we can think of these estimating functions as estimators of a risk function, which by definition is the expectation of a loss function. For instance if the loss is $\phi(\theta, Y) = -\log g^\theta(Y)$, $\Phi_{\bar{O}_n}(\theta)$ estimates the risk $E_*\{-\log g^\theta(Y)\}$; this risk is the cross entropy of g^θ relative to f^* (see section 3.1). Maximum likelihood estimators (MLE) are derived this way (Akaike, 1973; Commenges, 2009). A more general class of estimators of this form is that of M-estimators. Under some conditions given in Van der Vaart (2000), $\hat{\theta}$ converges in probability toward $\theta_0 = \operatorname{argmin}_{\theta} \Phi_\infty(\theta)$. Several estimating loss functions depend on θ only through g^θ but some may depend on θ directly, for instance through penalty terms.

Let us consider another loss function $\psi(g^\theta, Y)$ which defines a risk function $R^\psi(g^\theta) = E_*\{\psi(g^\theta, Y)\}$; we shall assume that $\psi(g^\theta, y)$ is a twice differentiable function in θ for all values of y . This risk function may serve for assessing estimators, but since estimators are random, the assessment risk is an expected risk: $ER^\psi(g^{\hat{\theta}}) = E_*\{\psi(g^{\hat{\theta}}, Y)\}$. Estimators with small assessment risks are preferred. The problem is to estimate the assessment risk (without knowing the true density f^*). This will be easier if the loss itself does not involve f^* ; this is the case for instance of $\psi(\theta, Y_i) = -\log g^\theta(Y_i)$ but not of the integrated squared error $ISE(g^\theta) = \int \{g^\theta(u) - f^*(u)\}^2 du$. The expectation however is with respect to the true distribution, so the risk cannot be computed exactly but has to be estimated. If another sample $\bar{O}'_n = (Y'_i, i = 1, \dots, n)$ iid with respect to \bar{O}_n were available, a natural estimator of the risk would be $\tilde{ER}^\psi(g^{\hat{\theta}}) = n^{-1} \sum_{i=1}^n \psi(g^{\hat{\theta}}, Y'_i)$. This is an unbiased estimator of the risk. This is often used by practitioners who split their original sample in a training and a validation sample. However this practice leads to a loss of efficiency since only half of the data are used for computing the estimator $g^{\hat{\theta}}$ and half of the data also for estimating its assessment risk. Crossvalidation estimators make a more efficient use of the information.

In particular the leave-one-out crossvalidation criterion is:

$$\text{CV}(g^{\hat{\theta}}) = n^{-1} \sum_{i=1}^n \psi(g^{\hat{\theta}_{-i}}, Y_i),$$

where $\hat{\theta}_{-i} = \text{argmin } \Phi_{\bar{\mathcal{O}}_{n|i}}$ and $\Phi_{\bar{\mathcal{O}}_{n|i}} = \frac{1}{n-1} \sum_{j \neq i}^n \phi(\theta, Y_j)$. CV does nearly as well as if another sample $\bar{\mathcal{O}}'_n$ were available. Indeed it can immediately be seen that $E\{\text{CV}(g^{\hat{\theta}})\} = E\text{R}^{\psi}(g^{\hat{\theta}_{n-1}})$. We shall see its asymptotic distribution in section 4. For comparing two estimators the difference of risks is relevant. This can of course be estimated by the difference of the estimated risks whose asymptotic distribution will also be given in section 4.

2.2 The universal approximate crossvalidation criterion

The leave-one-out crossvalidation criterion may be computationally demanding since it is necessary to run the maximization algorithm n times for finding the $\hat{\theta}_{-i}, i = 1, \dots, n$. For this reason an approximate formula is very useful. This is possible because $\hat{\theta}_{-i}$ is close to $\hat{\theta}$. Indeed, under mild conditions given in Van der Vaart (2000), $\sqrt{n}(\hat{\theta} - \theta_0)$ has an asymptotic normal distribution; since this is also true for $\hat{\theta}_{-i}$, this implies that $\hat{\theta}_{-i} - \hat{\theta} = O_p(n^{-1/2})$. A Taylor expansion of $\frac{\partial \Phi_{\bar{\mathcal{O}}_{n|i}}}{\partial \theta} |_{\hat{\theta}_{-i}}$ around $\hat{\theta}$ yields:

$$\hat{\theta}_{-i} - \hat{\theta} = -H_{\Phi_{\bar{\mathcal{O}}_{n|i}}}^{-1} \frac{\partial \Phi_{\bar{\mathcal{O}}_{n|i}}}{\partial \theta} |_{\hat{\theta}} + R_n,$$

where $H_{\Phi_{\bar{\mathcal{O}}_{n|i}}} = \frac{\partial^2 \Phi_{\bar{\mathcal{O}}_{n|i}}}{\partial \theta^2} |_{\hat{\theta}}$, and R_n is a quadratic form of $\hat{\theta}_{-i} - \hat{\theta}$ involving second and third derivatives of $\Phi_{\bar{\mathcal{O}}_{n|i}}$ taken in $\tilde{\theta}$ so that $\|\tilde{\theta}_n - \hat{\theta}\| \leq \|\hat{\theta}_{-i} - \hat{\theta}\|$. Thus $\|\tilde{\theta}_n - \hat{\theta}\|$ is also an $O_p(n^{-1/2})$. In virtue of the strong law of large numbers $\Phi_{\bar{\mathcal{O}}_{n|i}}$ and its derivatives converge almost surely toward their expectations taken at θ_0 . By definition of $\Phi_{\bar{\mathcal{O}}_n}(\theta)$ we have the relation:

$$n\Phi_{\bar{\mathcal{O}}_n}(\theta) = (n-1)\Phi_{\bar{\mathcal{O}}_{n|i}}(\theta) + \phi(\theta, Y_i). \quad (1)$$

Taking derivatives of the terms of this equation and taking the values at $\hat{\theta}$ we find $0 = (n-1) \frac{\partial \Phi_{\bar{\mathcal{O}}_{n|i}}}{\partial \theta} |_{\hat{\theta}} + \frac{\partial \phi(\theta, Y_i)}{\partial \theta} |_{\hat{\theta}}$ and we obtain that $\frac{\partial \Phi_{\bar{\mathcal{O}}_{n|i}}}{\partial \theta} |_{\hat{\theta}} = -\hat{d}_i$ where

$$\hat{d}_i = \frac{1}{n-1} \frac{\partial \phi(\theta, Y_i)}{\partial \theta} |_{\hat{\theta}}. \quad (2)$$

Note that on the right hand of the equation we could use a multiplicative factor $\frac{1}{n}$ instead of $\frac{1}{n-1}$ since the difference between the two is a $O(n^{-2})$. Hence we have:

$$\hat{\theta}_{-i} - \hat{\theta} = H_{\Phi_{\bar{\mathcal{O}}_{n|i}}}^{-1} \hat{d}_i + R_n, \quad (3)$$

Note that this implies that $\hat{\theta}_{-i} - \hat{\theta} = O_p(n^{-1})$ because $H_{\Phi_{\bar{\mathcal{O}}_{n|i}}} = O_p(1)$ and both \hat{d}_i and R_n are $O_p(n^{-1})$. But this in turn implies that R_n is in fact an $O_p(n^{-2})$. By twice

derivating (1) we obtain: $H_{\Phi_{\mathcal{O}_n}} = \frac{n-1}{n} H_{\Phi_{\mathcal{O}_{n|i}}} + \frac{1}{n} H_{\Phi_{\mathcal{O}_i}}$ where $H_{\Phi_{\mathcal{O}_i}} = \frac{\partial^2 \Phi_{\mathcal{O}_i}}{\partial \theta^2} |_{\hat{\theta}}$; since the last term is an $O_p(n^{-1})$, we can replace $H_{\Phi_{\mathcal{O}_{n|i}}}$ by $H_{\Phi_{\mathcal{O}_n}} = \frac{\partial^2 \Phi_{\mathcal{O}_n}}{\partial \theta^2} |_{\hat{\theta}}$ in (3) and obtain:

$$\hat{\theta}_{-i} - \hat{\theta} = H_{\Phi_{\mathcal{O}_n}}^{-1} \hat{d}_i + O_p(n^{-2}). \quad (4)$$

Developing now the assessment loss function for $\hat{\theta}_{-i}$ around $\hat{\theta}$ yields:

$$\psi(g^{\hat{\theta}_{-i}}, Y_i) = \psi(g^{\hat{\theta}}, Y_i) + (\hat{\theta}_{-i} - \hat{\theta})^T \hat{v}_i + O_p(n^{-2}),$$

where

$$\hat{v}_i = \frac{\partial \psi(g^{\hat{\theta}}, Y_i)}{\partial \theta} |_{\hat{\theta}}. \quad (5)$$

Replacing in this equation $\hat{\theta}_{-i} - \hat{\theta}$ by its approximation in (4) we obtain: $\psi(g^{\hat{\theta}_{-i}}, Y_i) = \psi(g^{\hat{\theta}}, Y_i) + \hat{d}_i^T H_{\Phi_{\mathcal{O}_n}}^{-1} \hat{v}_i + O_p(n^{-2})$. Taking the mean of the left terms of these equations yields $\text{CV}(g^{\hat{\theta}})$ so that we have:

$$\text{CV}(g^{\hat{\theta}}) = \Psi(g^{\hat{\theta}}) + \text{Trace}(H_{\Phi_{\mathcal{O}_n}}^{-1} K) + O_p(n^{-2}), \quad (6)$$

where $\Psi(g^{\hat{\theta}}) = n^{-1} \sum_{i=1}^n \psi(g^{\hat{\theta}}, Y_i)$ and $K = n^{-1} \sum_{i=1}^n \hat{v}_i \hat{d}_i^T$. We define UACV as:

$$\text{UACV}(g^{\hat{\theta}}) = \Psi(g^{\hat{\theta}}) + \text{Trace}(H_{\Phi_{\mathcal{O}_n}}^{-1} K). \quad (7)$$

2.3 Extension to incompletely observed data

More generally the available observation is $\bar{\mathcal{O}}_n = (\mathcal{O}_i, i = 1, \dots, n)$ where \mathcal{O}_i are sigma-fields containing observation of events generated by Y_i . Classical cases are that of right-censored observations common in survival analysis, left-censored observations common in biological measurements with detection limits, interval-censored data common in observations from cohort studies. Then risks functions of the type $\phi(\theta, Y_i)$ and $\psi(g^{\hat{\theta}}, Y_i)$ cannot be computed from observations (they are not $\bar{\mathcal{O}}_n$ -measurable). The theory above does not apply and attempts to estimate such functions require that the model is well-specified, a very strong assumption in this context. A solution is to consider loss functions of the form $\phi(\theta, \mathcal{O}_i)$ and $\psi(g^{\hat{\theta}}, \mathcal{O}_i)$, to which the above results apply. The most natural choice is the log-likelihood. Such an approach is in fact heuristically taken when using AIC with censored data. Some theory and simulations have been given in Commenges et al. (2007), Liqueur and Commenges (2011) and Commenges et al. (2012).

3 Particular cases of UACV

3.1 Maximum likelihood estimators and information risk: AIC

Suppose we take: $\phi(\theta, Y_i) = \psi(g^\theta, Y_i) = -\log g^\theta(Y_i)$. Then, the estimating function is minus the loglikelihood which estimates the estimating risk, here the cross-entropy (Cover and Thomas, 1991) of g^θ with respect to the true density f^* : $E_*\{-\log g^\theta(Y)\} = H(f^*) + \text{KL}(g^\theta; f^*)$, where $H(f^*) = -E_*\{\log f^*(Y)\}$ is the entropy of f^* and $\text{KL}(g^\theta; f^*) = E_*\{\log \frac{f^*(Y)}{g^\theta(Y)}\}$ the Kullback-Leibler divergence of g^θ relative to f^* . As for the assessment risk, this is the expected cross entropy:

$$\text{ECE}(g^{\hat{\theta}}) = E_*[E_*\{-\log g^{\hat{\theta}}(Y)|\bar{\mathcal{O}}_n\}] = H(f^*) + \text{EKL}(g^{\hat{\theta}}; f^*), \quad (8)$$

where $\text{EKL}(g^{\hat{\theta}}; f^*) = E_*\{\log \frac{f^*(Y)}{g^{\hat{\theta}}(Y)}\}$. In that case the loss functions for estimating and assessment are the same. The estimating and assessment risks are nearly the same; there is however a dissymmetry in that the estimating risk is a cross entropy while, because $g^{\hat{\theta}}$ is random, the assessment risk is an expected cross-entropy.

In that case the leading term of (7) is minus the maximized (normalized) loglikelihood. We have also that \hat{v}_i is the individual score and $\hat{d}_i = \frac{1}{n-1} \hat{v}_i$ so that UACV is identical to a normalized version of Takeuchi information criterion (TIC). If the model is well specified K tends in probability toward I_L , where I_L is the individual information matrix. The Hessian $H_{\Phi_{\mathcal{O}_n}}^{-1}$ also tends toward I_L so that the correction term tends toward p , the number of parameters. Thus, if the model is not too badly specified, TIC is approximately equal to AIC. We have $\text{UACV} = \frac{1}{2n} \text{TIC} \approx \frac{1}{2n} \text{AIC}$, and this estimates the *expected* cross-entropy of the estimator.

3.2 Maximum a posteriori and maximum penalized likelihood estimators

Estimators can be obtained by maximizing the maximum a posteriori (MAP) density of θ : $\sum_{i=1}^n \log g^\theta(Y_i) - J(\theta)$, where $J(\theta)$ is minus the log of the prior density of θ . This corresponds to a loss function $\phi(\theta, Y) = -\log g^\theta(Y) + n^{-1} J(\theta)$. Here the loss function depends on n but since the MAP estimators also converge (toward the same limit as the MLE), the above results still hold so that UACV can be applied to the MAP estimators. The assessment risk can be the expected cross entropy. Thus, in that case the estimating risk and the assessment risk are clearly different.

In the case where at least one of the parameters is a function, penalized likelihood can be applied for obtaining a smooth estimator (O’Sullivan, 1986). The function is approximated on a spline basis so that the problem reduces to a parametric one with a large number of parameters. This leads to $\phi(\theta, Y_i) = -\log g(Y_i; \theta) + J(\theta, n)$, where $J(\theta, n)$ is a penalty, which involves for instance the norm of the second derivative of the function and which may depend on n . $J(\theta, n)$ may depend on n in such a way that the estimator converges. The assessment risk can be the expected cross entropy as in section 3.1. In that case the leading term in (7) is minus the (normalized) loglikelihood taken at the value of the penalized likelihood estimator, $\hat{\theta}$. We still have that \hat{v}_i is the individual score (although taken at the value of the penalized likelihood estimator), and $\hat{d}_i = \frac{1}{n-1}(\hat{v}_i + \frac{\partial J(\theta, n)}{\partial \theta}|_{\hat{\theta}})$. We have that $H_{\Phi_{\odot_n}} = -\frac{\partial^2 L_{\odot_n}}{\partial \theta^2}|_{\hat{\theta}} + \frac{\partial^2 J(\theta, n)}{\partial \theta^2}|_{\hat{\theta}}$. In the case where $\frac{\partial^2 J(\theta, n)}{\partial \theta^2}$ is a definite positive matrix, we have $H_{\Phi_{\odot_n}} > -\frac{\partial^2 L_{\odot_n}}{\partial \theta^2}|_{\hat{\theta}}$ and $H_{\Phi_{\odot_n}}^{-1} < -\left(\frac{\partial^2 L_{\odot_n}}{\partial \theta^2}|_{\hat{\theta}}\right)^{-1}$ so that the correction term is smaller than for MLE. The correction term is often interpreted as the equivalent number of parameters (Commenges et al., 2007).

3.3 Hierarchical likelihood estimators

Let us consider the following model: conditionally on b_i , Y_i has a density $g_{Y|b}(\cdot; \theta, b_i)$, where b_i are random effects (or parameters). The (Y_i, b_i) are iid, where Y_i is multivariate of dimension n_i . The b_i are assumed to have density $g_b(\cdot; \tau)$ with zero expectation, where τ is a vector of parameters. Typically Y_i is (at least partially) observed while b_i is not.

Estimators of both θ and $\mathbf{b} = (b_1, \dots, b_n)$ are defined as maximizing the h-loglikelihood which can be written: $\text{HL}(\theta, \mathbf{b}, \tau) = \frac{1}{n} \sum_{i=1}^n \text{hl}(Y_i; \theta, b_i, \tau)$ with $\text{hl}(Y_i; \theta, b_i, \tau) = g_Y(Y_i; \theta, b_i) + \log g_b(b_i; \tau)$.

Commenges et al. (2011) showed (via a profile likelihood argument) that the hierarchical likelihood estimators for θ are M-estimators: $\text{PHL}_n(\theta) = \text{HL}(\theta, \hat{\mathbf{b}}(\theta)) = \frac{1}{n} \sum_{i=1}^n \text{hl}(Y_i; \theta, \hat{b}_i(\theta))$. Thus the estimating loss is $\text{hl}(Y_i; \theta, \hat{b}_i(\theta))$. Estimators may be compared using UACV. If we use the expected cross entropy for assessing the estimators, it is necessary to compute the marginal likelihood in $\hat{\theta}$ and the values of the individual scores. Of course, this is precisely what we attempted to avoid by using hierarchical likelihood; we have however to do it only once (rather than repeatedly within an iterative algorithm), so it may be computationally feasible. For the correcting term we have also to compute the Hessian of the hierarchical likelihood and the individual scores.

3.4 Restricted AIC

Liquet and Commenges (2011) have proposed a modification of AIC and LCV for assessing estimators based on information $\bar{\mathcal{O}}_n$ while the assessment risk is based on a smaller information $\mathcal{O}'_{n+1} \subset \mathcal{O}_{n+1}$. More specifically, the estimator is based on the sample $\bar{\mathcal{O}}_n = (Y_i, i = 1, \dots, n)$ but the assessment risk is based on a random variable Z which is a coarsened version of Y . For instance Z is a dichotomization of Y : $Z = 1_{Y > l}$. For this case, the restricted AIC (RAIC) was derived by both direct approximation of the risk and by approximation of the LCV. RAIC is clearly a particular case of UACV for the case: $\phi(\theta, Y_i) = -\log g^\theta(Y_i)$ and $\psi(g^\theta(Y_i)) = -\log g^\theta(Z_i)$.

3.5 Case where the estimators are assessed by CRPS

Gneiting and Raftery (2007) studied scoring rules and particularly the continuous rank probability score (CRPS). Its inverse can be used as a loss function and is defined as:

$$\text{CRPS}^*(G(\cdot, \theta), Y) = \int_{-\infty}^{+\infty} \{G(u, \theta) - 1_{u \geq Y}\}^2 du,$$

where $G(\cdot, \theta)$ is the cumulative distribution function (cdf) of a distribution in the model. The risk is a Cramer-von Mises type distance : $d(G, G^*) = \int \{G(u) - G^*(u)\}^2 du$. It may be interesting to assess MLE's using this assessment risk. In UACV, the leading term is $n^{-1} \sum_{i=1}^n \text{CRPS}^*(G(\cdot, \hat{\theta}), Y_i)$; for the correcting term, $H_{\Phi_{\bar{\mathcal{O}}_n}}$ is the Hessian of the loglikelihood and K must be computed with $\hat{v}_i = \frac{\partial \psi}{\partial \theta} |_{\hat{\theta}} = 2 \int_{-\infty}^{+\infty} \{G(u, \hat{\theta}) - 1_{u \geq Y}\} \frac{\partial G(u, \theta)}{\partial \theta} |_{\hat{\theta}} du$.

3.6 Estimators based on continuous approximation of categorical data

Assume Y is an ordered categorical variable taking values $l = 0, 1, \dots, L$. Here for simplicity we consider that Y is univariate. Several models are available for this type of variables. Threshold link models assume that $Y_i = l$ if a latent variable Λ_i takes values in the interval (c_l, c_{l+1}) for $l = 1, \dots, L$, with $c_{L+1} = +\infty$ and $Y_i = 0$ if $\Lambda_i < c_1$:

$$Y_i = \sum_{l=1}^{L+1} 1_{\{\Lambda_i \in (c_l, c_{l+1})\}} l. \quad (9)$$

Λ_i itself can be modeled as a noisy linear form of explanatory variables $\Lambda_i = \beta x_i + \varepsilon_i$, with ε_i having a normal distribution of mean zero and variance σ^2 , and where x_i are explanatory variables. The parameters are $\theta = (c_1, \dots, c_L, \beta, \sigma)$. An estimator of the distribution can be obtained by maximum likelihood leading to define $g^{\hat{\theta}}$. For assessing the estimator it

is natural to use $\text{ECE}(g^{\hat{\theta}})$ (defined in 8). Note that since Y is discrete the densities are defined with respect to a counting measure that is, $g^{\hat{\theta}}(l)$ defines the probability that $Y = l$.

One may also make a continuous approximation which leads to simpler computations and may be more parsimonious, especially if Y is multivariate as in the illustration of section 6. In this approach we consider the model $Y_i = \beta x_i + \varepsilon_i$. Maximizing the likelihood of this model for observations of Y_i leads to a probability measure specified by the density $h_c^{\hat{\gamma}}$. This is however a density relative to Lebesgue measure. This probability measure gives zero probabilities to $\{Y_i = l\}$ for all l , and this yields infinite value for ECE (meaning strong rejection of this estimator). However from h_c a natural estimator of f^* can be constructed by gathering at l the mass around l : $h^{\hat{\gamma}}(l) = \int_{l-1/2}^{l+1/2} h_c^{\hat{\gamma}}(u) du$, for $l = 1, \dots, L-1$, and $h^{\hat{\gamma}}(0) = \int_{-\infty}^{1/2} h_c^{\hat{\gamma}}(u) du$, $h^{\hat{\gamma}}(L) = \int_{L-1/2}^{+\infty} h_c^{\hat{\gamma}}(u) du$. UACV can be computed for this estimator for estimating its ECE. The first term of $\text{UACV}(h^{\hat{\gamma}})$ can be interpreted as the loglikelihood obtained by this estimator with respect to the counting measure. For the correcting term we need the Hessian of the loglikelihood of $h_c^{\hat{\gamma}}$ and we have to compute $\hat{v}_i = \frac{\partial \psi(h^{\hat{\gamma}}, Y_i)}{\partial \gamma} \Big|_{\hat{\theta}}$. For instance if $Y_i = l$ for $l = 1 \dots, L-1$ we have

$$\hat{v}_i = - \frac{\int_{l-1/2}^{l+1/2} \frac{\partial h_c^{\hat{\gamma}}}{\partial \gamma}(u) du}{\int_{l-1/2}^{l+1/2} h_c^{\hat{\gamma}}(u) du}.$$

Since the denominator is the probability under $h_c^{\hat{\gamma}}$ that $Y \in (l-1/2, l+1/2)$, \hat{v}_i can be interpreted as the conditional expectation (under $h_c^{\hat{\gamma}}$) of the individual score. Thus if $h_c^{\hat{\gamma}}$ does not vary much on $(l-1/2, l+1/2)$, \hat{v}_i is close to $-(n-1)\hat{d}_i$. Using the same arguments as in section 3.1 we obtain that UACV is close to correcting by the number of parameters as in AIC, a criterion that we call AIC_d . This is what Proust-Lima et al. (2011) proposed, and this is likely to be a good approximation if the number of modalities of Y is large.

4 Asymptotic distribution and tracking interval

Commenges et al. (2008) using results of Vuong (1989) studied the asymptotic distribution of a normalized difference of AIC as an estimator of a difference of Kullback-Leibler risks. Here similar arguments are applied to study the asymptotic distribution of UACV and a difference of UACV. Since UACV and CV differ by an $O_p(n^{-2})$, this will also give the asymptotic distribution of CV and a difference of CV. By the continuous mapping theorem, the asymptotic distribution of $\text{UACV}(g^{\hat{\theta}})$ is the same as that of $\Psi(g^{\theta_0})$. Since the latter

quantity is a mean, it immediately follows by the central limit theorem that:

$$n^{1/2}\{\text{UACV}(g^{\hat{\theta}}) - R^{\psi}(g^{\theta_0})\} \xrightarrow{D} \mathcal{N}(0, \kappa_*^2), \quad (10)$$

where $\kappa_*^2 = \text{var}_* \psi(g^{\theta_0}, Y)$. We can also write:

$$n^{1/2}\{\text{UACV}(g^{\hat{\theta}}) - \text{ER}^{\psi}(g^{\hat{\theta}})\} \xrightarrow{D} \mathcal{N}(0, \kappa_*^2), \quad (11)$$

and κ_*^2 can be estimated by the empirical variance of $\psi(g^{\hat{\theta}}, Y_i)$, $i = 1, \dots, n$.

Given two estimators, $g^{\hat{\theta}}$ and $h^{\hat{\gamma}}$, it is interesting to estimate the difference of their risks: $\Delta^{\psi}(g^{\hat{\theta}}, h^{\hat{\gamma}}) = \text{ER}^{\psi}(g^{\hat{\theta}}) - \text{ER}^{\psi}(h^{\hat{\gamma}})$. The obvious estimator is: $D_{\text{UACV}}(g^{\hat{\theta}}, h^{\hat{\gamma}}) = \text{UACV}(g^{\hat{\theta}}) - \text{UACV}(h^{\hat{\gamma}})$. We focus on the case where $g^{\beta_0} \neq h^{\gamma_0}$. We obtain in that case using the same arguments as above:

$$n^{1/2}\{D_{\text{UACV}}(g^{\hat{\beta}_n}, h^{\hat{\gamma}_n}) - \Delta(g^{\hat{\beta}_n}, h^{\hat{\gamma}_n})\} \xrightarrow{D} \mathcal{N}(0, \omega_*^2), \quad (12)$$

where $\omega_*^2 = \text{var} \{ \psi(g^{\theta_0}, Y) - \psi(h^{\gamma_0}, Y) \}$, and this can be estimated by the empirical variance of $\{ \psi(g^{\hat{\theta}}, Y_i) - \psi(h^{\hat{\gamma}}, Y_i) \}$.

From this we can compute the so called tracking interval (A_n, B_n) , where $A_n = D_{\text{UACV}}(g^{\hat{\beta}_n}, h^{\hat{\gamma}_n}) - z_{\alpha/2} n^{-1/2} \hat{\omega}_n$ and $B_n = D_{\text{UACV}}(g^{\hat{\beta}_n}, h^{\hat{\gamma}_n}) + z_{\alpha/2} n^{-1/2} \hat{\omega}_n$, where z_u is the u^{th} quantile of the standard normal variable.

Note that ω_* is in general much lower than κ_* . This has been shown by Commenges et al. (2007) for the cross entropy risk and comes from the fact that $\psi(g^{\hat{\theta}}, Y_i)$ and $\psi(h^{\hat{\gamma}}, Y_i)$ are often positively correlated.

5 Simulation: choice of estimators for ordered categorical data

5.1 Design

We conducted a simulation study to illustrate the use of UACV for comparing estimators derived from threshold link models and estimators obtained by a continuous approximation in the case of ordered categorical data (see section 3.6). The aim was to assess the performance of UACV as an estimator of ECE, and to compare it to the normalized naive AIC criterion (noted AIC) and the normalized AIC criterion computed on the counting measure (noted AIC_d). Performances of these criteria were studied in the case where the number of modalities $(L + 1)$ of the response variable Y is small (section 5.2), and where L is large (section 5.3).

True distributions

For all the simulations, the data came from a threshold model specified by:

$$\begin{cases} Y_i = l & \text{if } \Lambda_i \in [c_l; c_{l+1}) \text{ for } l = 1, \dots, L \\ Y_i = 0 & \text{if } \Lambda_i < c_1 \text{ and } c_{L+1} = +\infty \\ \Lambda_i = \beta_0 + \beta_1 X_i^1 + \beta_2 X_i^2 + \varepsilon_i & i = 1, \dots, n \end{cases}$$

where ε_i had a normal distribution of mean zero and variance σ^2 and the two explanatory variables X^1 and X^2 came from a standard normal distribution. In order to not disadvantage the linear continuous approximation compared to the threshold model, the parameters c_1, \dots, c_L were chosen as the solution of the following equations:

$$\begin{cases} P(Y_i < c_1) = P(Y_i > c_L) \\ P(Y_i < c_1) = P(c_1 < Y_i < c_2), \\ c_{i+1} = c_i + m \text{ with } m = (c_L - c_1)/(L - 1) \end{cases}$$

The different models

For each generated sample, we fitted the threshold model as previously defined, and a linear model assuming a linear continuous approximation of the response variable Y , $Y_i = \beta_0 + \beta_1 X_i^1 + \beta_2 X_i^2 + \varepsilon_i$. Both models were estimated by maximum likelihood.

For all simulations, $N = 1000$ samples were generated. The true criterion ECE (available by simulation) was computed by Monte-carlo approach.

5.2 Small number of modalities ($L = 4$)

We consider here the case where the number of modalities of Y is relatively small ($L + 1 = 5$). In this simulation, we fixed $\beta_0 = 1$, $\beta_1 = -2.1$, $\beta_2 = -3.7$ and $\sigma^2 = 4$. In Table 1, we present, for different sample sizes n , the results for the different empirical criteria AIC, AIC_d and UACV which can be compared to the true criterion ECE. For any sample size, the threshold model provided a better ECE than the linear model (positive difference). It appeared that UACV had a small bias for all the sample sizes (of order 10^{-3}). The two others criteria AIC and AIC_d were also in favor of a threshold model. However, as expected, the naive normalized AIC failed to estimate the true ECE risk due to the wrong probability measure (Lebesgue measure instead of a counting measure). We can note that the criterion AIC_d estimated ECE relatively well, with a bias around 10^{-2} and 10^{-3} .

Table 1: Performance of the criteria for a small number of modalities ($L + 1 = 5$) and different sample sizes. Mean over 1000 replications of the difference of criteria UACV, AIC_d , AIC. ECE is the true risk. Biases of the criteria as estimator of ECE

	ECE	UACV	AIC_d	AIC	Bias UACV	Bias AIC_d	Bias AIC
n=300							
Linear	1.2650	1.2661	1.2702	1.5998	0.0011	0.0053	0.3348
Threshold	0.9874	0.9835	0.9836	0.9836	-0.0039	-0.0038	-0.0038
Difference	0.2775	0.2826	0.2866	0.6162	0.0051	0.0091	0.3387
n=500							
Linear	1.2594	1.2635	1.2660	1.6468	0.0041	0.0066	0.3875
Threshold	0.9753	0.9810	0.9810	0.9810	0.0057	0.0057	0.0057
Difference	0.2840	0.2825	0.2849	0.6658	-0.0016	0.0009	0.3818
n=3000							
Linear	1.2612	1.2617	1.2621	1.6325	0.0005	0.0009	0.3713
Threshold	0.9763	0.9749	0.9749	0.9749	-0.0014	-0.0014	-0.0014
Difference	0.2850	0.2868	0.2873	0.6577	0.0019	0.0023	0.3727

5.3 Large number of modalities

We consider here the case where the number of modalities of Y is relatively large ($L + 1 = 20$). In this simulation, we fixed $\beta_0 = 1$, $\beta_1 = -0.3$, $\beta_2 = -1.7$ and $\sigma^2 = 4$. The results of this simulation are presented in Table 2. For any sample size, the linear model provided a better ECE than the threshold model (negative difference). It appeared that UACV had a small bias for all the sample sizes (of order 10^{-3} and 10^{-4}). The AIC_d criterion gave similar results as the UACV criterion while the AIC criterion failed to find the best estimator (positive difference).

6 Illustration on the choice of estimators for psychometric tests

In epidemiological studies, cognition is measured by psychometric tests which usually consist in the sum of items measuring one or several cognitive domains. A common example is the Mini Mental State Examination (MMSE) score (Folstein et al., 1975), computed as the

Table 2: Performance of the criteria for a large number of modalities ($L + 1 = 20$) and different sample sizes. Mean over 1000 replications of the difference of criteria UACV, AIC_d , AIC. ECE is the true risk. Biases of the criteria as estimator of ECE

	ECE	UACV	AIC_d	AIC	Bias UACV	Bias AIC_d	Bias AIC
n=300							
Linear	2.6794	2.6792	2.6796	2.8076	-0.0002	0.0002	0.1282
Threshold	2.7104	2.7070	2.7069	2.7069	-0.0035	-0.0036	-0.0036
Difference	-0.0311	-0.0278	-0.0273	0.1007	0.0033	0.0038	0.1317
n=500							
Linear	2.6765	2.6740	2.6742	2.7117	-0.0026	-0.0024	0.0352
Threshold	2.6933	2.6899	2.6899	2.6899	-0.0033	-0.0034	-0.0034
Difference	-0.0167	-0.0160	-0.0157	0.0218	0.0008	0.0010	0.0385
n=3000							
Linear	2.6736	2.6716	2.6717	2.7107	-0.0020	-0.0019	0.0371
Threshold	2.6746	2.6725	2.6725	2.6725	-0.0021	-0.0021	-0.0021
Difference	-0.0010	-0.0009	-0.0008	0.0382	0.0001	0.0001	0.0391

sum of 30 binary items (grouped in 20 independent items) evaluating memory, calculation, orientation in space and time, language, and word recognition; for this reason it is called a "sumscore" and ranges from 0 to 30. Although in essence psychometric tests are ordered categorical data, they are most often analyzed as continuous data. Indeed, they usually have a large number of different levels and, especially in longitudinal studies, models for categorical data are numerically complex. Recently, Proust-Lima et al. (2011) defined a latent process mixed model to analyse repeated measures of discrete outcomes involving either a threshold model or an approximation of it using continuous parameterized increasing functions. Comparison of models assuming either categorical data (using the threshold model) or continuous data (using continuous functions) was done with an AIC_d , computed with respect to the counting measure. In this illustration, we use UACV to compare such latent process mixed models assuming either continuous or ordered categorical data when applied on the repeated measures of the MMSE and its calculation subscore in a large sample from a French prospective cohort study.

6.1 Latent process mixed models

In brief, the latent process mixed model assumes that a latent process $(\Lambda_i(t))_{t \geq 0}$ underlies the repeated measures of the observed variable $Y_i(t_{ij})$ for subject i ($i = 1, \dots, n$) and occasion j ($j = 1, \dots, n_i$). First, the latent process trajectory is defined by a standard linear mixed model (Laird and Ware, 1982): $\Lambda_i(t) = X_i(t)^T \beta + Z_i(t)^T b_i + \epsilon_i(t_{ij})$ for $t \geq 0$ where $X_i(t)$ and $Z_i(t)$ are distinct vectors of time-dependent covariates associated respectively with the vector of fixed effects β and the vector of random effects b_i ($b_i \sim MVN(\mu, D)$), and $\epsilon_i(t) \sim \mathcal{N}(0, \sigma^2)$. We further assume that b_{i0} , the first component of b_i that usually represents the random intercept, is $N(0, 1)$ for identifiability; except for the variance of b_{i0} , D is an unstructured variance matrix.

Then, a measurement model links the latent process with the observed repeated measures. For ordered categorical data, a standard threshold model as defined in (9) (section 3.6) for the univariate case is well adapted. For continuous data, the link has been modeled as $H(Y_i(t_{ij}); \eta) = \Lambda_i(t_{ij})$ where $H(\cdot; \eta)$ is a monotonic increasing transformation. Three families of such transformations are considered: (i) $H(y; \eta) = \frac{h(y; \eta_1, \eta_2) - \eta_3}{\eta_4}$ where $h(\cdot; \eta_1, \eta_2)$ is the Beta cdf with parameters (η_1, η_2) ; (ii) $H(y; \eta) = \eta_1 + \sum_{l=2}^{m+2} \eta_l B_l^I(y)$ where $(B_l^I)_{l=2, m+2}$ is a basis of quadratic I-splines with m nodes; (iii) $H(y; \eta) = \frac{y - \eta_1}{\eta_2}$ which gives the standard linear mixed model.

Latent process mixed models are estimated within the Maximum Likelihood framework using `lcm` function of `lcm` R package. When assuming continuous data, the log-likelihood can be computed analytically using the Jacobian of H . In contrast, when assuming ordered categorical data, an integration over the random-effects distribution intervenes in the likelihood computation which is approximated by a Gauss-Hermite quadrature. Further details can be found in Proust et al. (2006) and Proust-Lima et al. (2012).

UACV is computed from the loglikelihood Ψ obtained for the maximum likelihood estimators $\hat{\theta}$ with respect to the counting measure :

$$\begin{aligned} \Psi(\hat{\theta}) &= \sum_{i=1}^n \int_{b_i} \prod_{j=1}^{n_i} P(Y_{ij}|b_i) f_b(b_i) db_i \\ &= \sum_{i=1}^n \int_{b_i} \prod_{j=1}^{n_i} \prod_{l=0}^L (P(Y_{ij} = l|b_i))^{I_{Y_{ij}=l}} f_b(b_i) db_i \\ &= \sum_{i=1}^n \int_{b_i} \prod_{j=1}^{n_i} \prod_{l=0}^L (P(c_l \leq \Lambda_i(t_{ij}) + \epsilon_i(t_{ij}) < c_{l+1}|b_i))^{I_{Y_{ij}=l}} f_b(b_i) db_i \end{aligned} \quad (13)$$

where $\Lambda_i(t_{ij}) + \epsilon_i(t_{ij}) \sim N(X_i(t)^T \beta + Z_i(t)^T \mu, Z_i(t)^T D Z_i(t) + \sigma^2)$, $c_0 = -\infty$,




Figure 1: Distributions of MMSE Sumscore and MMSE Calculation Subscore in the PAQUID sample (n=2,914). Data were pooled from all available visits for a total of 10,846 observations.

$c_{L+1} = +\infty$, and either c_l ($l = 1, \dots, L$) are the estimated thresholds when a threshold model is considered, or $c_l = H(l - \frac{1}{2}, \hat{\eta})$ ($l = 1, \dots, L$) when monotonic increasing families of transformations are used.

6.2 Application: categorical psychometric tests

Data come from the French prospective cohort study PAQUID initiated in 1988 to study normal and pathological aging (Letenneur et al. (1994)). Subjects included in the cohort were 65 and older at initial visit and were followed up to 10 times with a visit at 1, 3, 5, 8, 10, 13, 15, 17 and 20 years after the initial visit. At each visit, a battery of psychometric tests including the MMSE was completed. In the present analysis, all the subjects free of dementia at the 1-year visit and who had at least one MMSE measure during the whole follow-up were included. Data from baseline were removed to avoid modeling the first-passing effect. A total of 2914 subjects were included with a median of 2 (Interquartile range=3-5) repeated measures of MMSE sumscore. The distributions of MMSE sumscore ranging from 0 to 30, and of its calculation subscore, ranging from 0 to 5, are displayed in Figure 2.

The trajectory of the latent process was modeled as an individual quadratic function of age with correlated random intercept, slope and quadratic slope ($Z_i(t) = (1, \text{age}_i, \text{age}_i^2)$), and an adjustment for binary covariates educational level (EL=1 if the subject graduated

from primary school) and gender (SEX=1 if the subject is a man) plus their interactions with age and quadratic age (so that $X_i(t) = Z_i(t) \otimes (1, EL_i, SEX_i)$). For MMSE sumscore, in addition to the threshold model, the linear, Beta cdf and I-splines (with 5 equidistant nodes) continuous link functions were considered. For calculation subscore, in addition to the threshold model, only the linear link was considered.

6.3 Results

Table 3 gives the assessment criteria for estimators based on the different models, and table 4 provides the differences in UACV or AIC and their 95% tracking interval. For the MMSE sumscore, the mixed model assuming the standard linear transformation yielded a clearly worse UACV than other models accounting for nonlinear relationships with the underlying latent process. The model involving a Beta cdf gave a similar risk as the one involving the less parsimonious I-splines transformation ($D_{UACV}(\text{Beta cdf}, \text{I-splines}) = -0.0070$, 95% Tracking interval: $[-0.0152, 0.0012]$). Finally, the mixed model considering a threshold link model, which is numerically demanding, gave the best fit but remained relatively close to the simpler ones assuming a Beta cdf ($D_{UACV}(\text{Beta cdf}, \text{Thresholds}) = 0.0200$, 95% Tracking interval: $[0.0097, 0.0303]$) or a I-splines transformation ($D_{UACV}(\text{I-splines}, \text{Thresholds}) = 0.0270$, 95% Tracking interval: $[0.0166, 0.0374]$). For the interpretation of these values Commenges et al. (2008) suggested to qualify values of order 10^{-1} , 10^{-2} and 10^{-3} as "large", "moderate" and "small" respectively; moreover for multivariate observations, it was suggested to divide by the total number of observations rather by the number of independent observations. With this correction (which amounts to divide the current values by a factor of $3.7 = 10846/2914$) the differences between the linear model and the other models can be qualified as "large", and the differences between the threshold model and both beta cdf and I-splines are between "moderate" and "small". Figure 2 displays the estimated link functions in (A) and the predicted mean trajectories of the latent process according to educational level in (B) from the models involving either a linear, a beta cdf, a I-splines or a threshold link function. The estimated link functions as well as the predicted trajectories of the latent process are very close when assuming either beta cdf, I-splines or a threshold link function but they greatly differ when assuming a linear link. This difference also shows up in the effects of covariates with associations distorted when not accounting for nonlinear transformations (as demonstrated in Proust-Lima et al. (2011)). For example in this application, a significant overall increase with age of the difference between educational levels is found when assuming a linear model ($p=0.011$ for interaction with age and

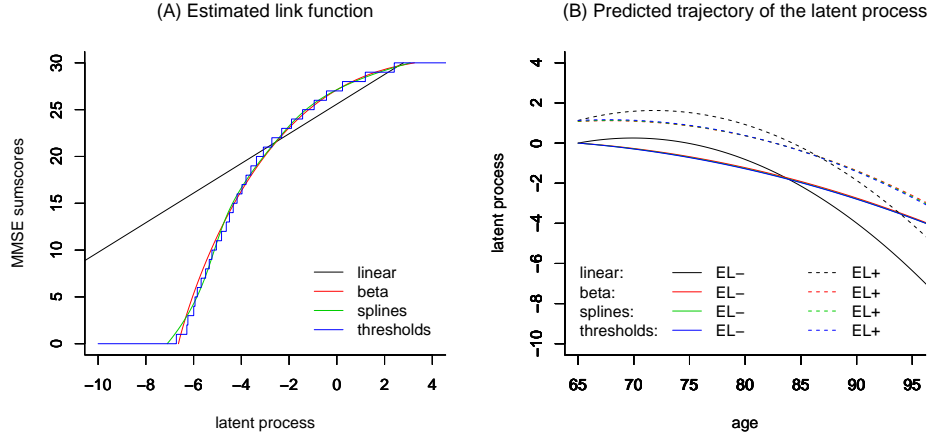


Figure 2: (A) estimated inverse link functions between MMSE sumscores and the underlying latent process, and (B) predicted trajectories of the latent process of a woman according to educational level (with EL+ and EL- for respectively validated or non-validated primary school diploma) in latent process mixed models assuming either linear, Beta cdf, I-splines or threshold link functions (PAQUID sample, $n=2,914$); the trajectories for the latter three transformations are indistinguishable

age squared) while it is a significant overall decrease with age which is found in the other models ($p\text{-value} < 0.001$ for interaction with age and age squared).

For the Calculation subscore also, the standard linear mixed model again gave a clearly higher risk than the mixed model assuming a threshold link model ($D_{UACV}(\text{linear}, \text{Thresholds}) = 0.4523$, 95% Tracking interval: $[0.4127, 0.4919]$).

7 Conclusion

We have proposed a universal approximate formula for leave-one out crossvalidation: it is universal in the sense that it applies to any couple of estimating and assessment risks which can be correctly estimated from the observations. This is in principle restricted to parametric models but extends to smooth semi- or non-parametric ones through penalized likelihood. The approximate formula not only allows fast computation, but also allows deriving the asymptotic distribution. Estimating this distribution is important since the variability of UACV, as that of all the criteria used for estimator choice, is large, even if the variability of a difference of UACV between two estimators is smaller.

Table 3: Number of parameters (p), individual log-likelihood ($\Phi(\hat{\theta})$), corresponding naive normalized AIC (AIC), log-likelihood computed with respect to the counting measure ($\Psi(\hat{\theta})$), corresponding AIC (AIC_d), and UACV for latent process mixed models involving different transformations H and applied on either the MMSE sumscore or its calculation subscore.

Transformation H	p	$\Phi(\hat{\theta})$	AIC	$\Psi(\hat{\theta})$	AIC_d	UACV
<i>MMSE</i>						
Linear	16	-8.7468	8.7523	-8.5231	8.5286	8.5361
Beta cdf [†]	18	-7.7514	7.7576	-7.7803	7.7865	7.7865
I-splines [‡]	21	-7.8249	7.8321	-7.7857	7.7929	7.7935
Thresholds	44	-7.7473	7.7624	-7.7473	7.7624	7.7665
<i>Calculation</i>						
Linear	16	-6.0057	6.0111	-4.8143	4.8215	4.8198
Thresholds	19	-4.3618	4.3683	-4.3618	4.3683	4.3692

[†] cdf for cumulative distribution function

[‡] Quadratic I-splines with 5 equidistant nodes located at 0, 7.5, 15, 22.5 and 30.

Table 4: Difference of AIC_d (D_{AIC_d}), difference of UACV and its 95% tracking interval between latent process mixed models involving different transformations H_1 and H_2 , and applied on either the MMSE sumscore or its calculation subscore.

Transformations H_1 / H_2	D_{AIC_d}	D_{UACV}	95% tracking interval
<i>MMSE</i>			
linear / Beta cdf [†]	0.7421	0.7495	[0.6619 ; 0.8372]
linear / I-splines [‡]	0.7357	0.7425	[0.6526 ; 0.8325]
Beta cdf [†] / I-splines [‡]	-0.0064	-0.0070	[-0.0152 ; 0.0012]
I-splines [‡] / thresholds	0.0306	0.0270	[0.0166 ; 0.0374]
Beta cdf [†] / thresholds	0.0241	0.0200	[0.0097 ; 0.0303]
linear / thresholds	0.7662	0.7696	[0.6784 ; 0.8607]
<i>Calculation</i>			
linear / thresholds	0.4515	0.4523	[0.4127 ; 0.4919]

[†] cdf for Cumulative Distribution Function

[‡] Quadratic I-splines with 5 equidistant nodes located at 0, 7.5, 15, 22.5 and 30.

In this paper, UACV has been applied to the issue of choice between estimators of the distribution of categorical data based on threshold models or on models based on a continuous approximation. It has been shown that the naive AIC can be misleading while a procedure called AIC_d (which had not been theoretically validated) yields results very close to UACV, even if the latter is slightly better. Both quantities can be computed in the `lcm` R package.

References

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In P. B. N. and C. F. (Eds.), *Proc. of the 2nd Int. Symp. on Information Theory*, pp. 267–281.
- Arlot, S. and A. Celisse (2010). A survey of cross-validation procedures for model selection. *Statist. Surv.* 4, 40–79.
- Commenges, D. (2009). Statistical models, likelihood, penalized likelihood and hierarchical likelihood. *Statistics Surveys* 3, 1–17.
- Commenges, D., D. Jolly, J. Drylewicz, H. Putter, and R. Thiébaud (2011). Inference in HIV dynamics models via hierarchical likelihood. *Computational Statistics & Data Analysis* 55, 446–456.
- Commenges, D., P. Joly, A. Gegout-Petit, and B. Liqueur (2007). Choice between semi-parametric estimators of markov and non-markov multi-state models from generally coarsened observations. *Scandinavian Journal of Statistics* 34, 33–52.
- Commenges, D., B. Liqueur, and C. Proust-Lima (2012). Choice of prognostic estimators in joint models by estimating differences of expected conditional kullback-leibler risks. *Biometrics* 68, in press.
- Commenges, D., A. Sayyareh, L. Letenneur, J. Guedj, and A. Bar-hen (2008). Estimating a difference of kullback-leibler risks using a normalized difference of aic. *Annals of Applied Statistics* 2(3), 1123–1142.
- Cover, T. and J. Thomas (1991). *Elements of information theory*, pp. 542. New York, NY: John Wiley and Sons.

- Folstein, M. F., S. E. Folstein, and P. R. McHugh (1975). "mini-Mental State". A practical method for grading the cognitive state of patients for the clinician. *Journal of Psychiatric Research* 12(3), 189–98.
- Gneiting, T. and A. Raftery (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477), 359–378.
- Golub, G., M. Heath, and G. Wahba (1979). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* 21(2), 215–223.
- Gu, C. and D. Xiang (2001). Cross-validating non-Gaussian data. *Journal of Computational and Graphical Statistics* 10(3), 581–591.
- Laird, N. and J. Ware (1982). Random-effects models for longitudinal data. *Biometrics* 38, 963–74.
- Letenneur, L., D. Commenges, J. F. Dartigues, and P. Barberger-Gateau (1994). Incidence of dementia and Alzheimer’s disease in elderly community residents of south-western France. *International Journal of Epidemiology* 23(6), 1256–61.
- Liquet, B. and D. Commenges (2011). Choice of estimators based on different observations: modified AIC and LCV criteria. *Scandinavian Journal of Statistics* 38, 268–287.
- O’Sullivan, F. (1986). A statistical perspective on ill-posed inverse problems. *Statistical Science* 1(4), 502–518.
- Proust, C., H. Jacqmin-Gadda, J. M. Taylor, J. Ganiayre, and D. Commenges (2006). A nonlinear model with latent process for cognitive evolution using multivariate longitudinal data. *Biometrics* 62(4), 1014–24.
- Proust-Lima, C., H. Amieva, and H. Jacqmin-Gadda (2012). Analysis of multivariate mixed longitudinal data: a flexible latent process approach. *submitted*.
- Proust-Lima, C., J. Dartigues, and H. Jacqmin-Gadda (2011). Misuse of the linear mixed model when evaluating risk factors of cognitive decline. *American Journal of Epidemiology* 174, 1077–1088.
- Stone, M. (1974). Cross-validatory choice and assessment of statistical predictions (with discussion). *Journal of the Royal Statistical Society B* 39, 111–147.

- Van Der Laan, M., S. Dudoit, and S. Keles (2004). Asymptotic optimality of likelihood-based cross-validation. *Statistical Applications in Genetics and Molecular Biology* 3(1), 1036.
- Van der Vaart, A. (2000). *Asymptotic Statistics*. Cambridge University Press.
- Vuong, Q. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica: Journal of the Econometric Society* 57(2), 307–333.
- Wahba, G. (1985). A comparison of GCV and GML for choosing the smoothing parameter in the generalized spline smoothing problem. *The Annals of Statistics*, 1378–1402.
- Xiang, D. and G. Wahba (1996). A generalized approximate cross validation for smoothing splines with non-Gaussian data. *Statistica Sinica* 6, 675–692.